Title: OLSRv2 Scalability

Contact:

Thomas Clausen, thomas.clausen@polytechnique.edu
Network Research Groupe of Ecole Polytechnique

The Optimized Link-State Routing Protocol (OLSR, and its successor OLSRv2) is *the IETF standardized routing protocol for Mobile Ad hoc NETworks (MANETs).* This includes MESH networks, etc.
A link state protocol, OLSR (RFC3626) prototyped several innovative techniques and principles, including the concept of Multi-Point Relays (MPRs) for Flooding Reduction and Topology Reduction, an advanced neighborhood discovery mechanism for detecting and tracking link-bi-directionality in wireless environments and across multiple wireless interfaces. Some of these techniques were later revised, adapted for OSPF and included in the OSPF extension for MANETs (see OSPF for MANETs).
OLSRv2 (RFC7181) and its constituent parts (RFC5148, RFC5444, RFC5497, RFC6130 RFC7181, RFC7182, RFC7183, RFC7184, RFC7185, RFC7186, RFC7187, RFC7188) were built on the 10+ years of experience with OLSR, and offers:

- High scalability in dense networks by way of Flooding Reduction and Topolgy Reduction
- A Modular plug-in security architecture
- Flexible, multi-metric routing
- Compressed (aggregated) addresses for significantly smaller control messages
- Built-in extensibility mechanisms, preserving forward and backwards compatibility
- Native and efficient IPv6 support
- Flexible and extensive gateway support, for attached networks as well as for global Internet connectivity.

For all its qualities, OLSRv2 is forming a single routing domain - which, despite claims of "high scalability" by way of flooding and topology reduction, all the same puts limits to the size of OLSRv2 deployments.

The Internet deals with scalability through "divide and conquer": divide "the world" into autonomous systems, each of which is routed independently by an IGP and has an AS number - then route between autonomous systems (as identified both by their AS numbers) using a different routing protocol (an EGP). Inside an AS, typically, the IGP operates by further subdividing into even smaller areas -- with an area limiting the flooding scope, as well as the amount of topology maintained in routers.

In other word, the Internet scales world wide because of (in part) the fact that it can be decomposed and organized into routing hierarchies -- this is, indeed, one of the "five dirty secrets that make the Net Work". This "decomposition and hierarchical organization" is statically configured, rarely changing, and accompanied with clever tricks such as assigning IPv6 addresses from within the same prefix to the same ASes/Areas.

In a MANET/MESH network, the same tricks are hard to apply: the spontaneous (ad hoc) nature means that static configuration, assigning of prefixes, and partitioning into domains is not evident and certainly not static ; addresses assigned to interfaces of devices that move will either have to change (challenging session continuity for user applications), or the relationship between IPv6 addresses and a physical location is not maintainable.

In short: bringing a MANET/MESH network to a world-wide scale is...a challenge. Fortunately, we have a couple of ideas that might allow OLSRv2 to scale. Testing - and if it works - making real those ideas, is what this project is about. Specifically, to:

1. Investigate the scalability limits of OLSRv2 as-is, comparing to its predecessor (OLSRv1), and using a selection of tools: network simulation, analytical models, etc.
2. As an initial approach to scalability, apply temporal partial topology techniques such as FSR and HSLS, as well as different flooding and topology reduction techniques (such as non-connected dominating sets).
3. In parallel to the two points above, pursue a literature survey on hierarchical routing in highly dynamic networks - not limited to (the few) approaches proposed for hierarchical OLSR
4. Participate, in the development and refinement of one innovative approach for establish dynamic routing hierarchies in OLSRv2, specification of the protocol extension and operation, and in devising and executing evaluation methods allowing to quantify both the scalability gains (if any) and possible inconveniences and induced complexities (if any).